

NBER WORKING PAPER SERIES

RELIGIOUS IDENTITY AND ECONOMIC BEHAVIOR

Daniel J. Benjamin
James J. Choi
Geoffrey W. Fisher

Working Paper 15925
<http://www.nber.org/papers/w15925>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2010

We thank Azim Shariff and Ara Norenzayan for sharing with us the religious identity priming instrument. For helpful comments and suggestions, we are grateful to Daniel Chen, Stefano DellaVigna, Kirabo Jackson, Sendhil Mullainathan, Ted O'Donoghue, participants at the NBER's Economics of Religion Conference and UC Santa Barbara's Behavioral/Experimental Economics Conference, and seminar participants at Maryland, Chapman, UNSW, ANU, Melbourne, Hebrew, Tel Aviv, Ben Gurion, and Bar-Ilan. We thank Mario Basora, Maria Bodiu, Kristen Cooper, Evan Buntrock, Jim Casteleiro, John Farragut, Isabel Fay, Rebecca Friedman, Joshua Funt, Arjun Gokhale, Jesse Gould, Rebecca Hausner, Ben Hebert, Liying Huang, Ahmed Jaber, Bige Kahraman, Anqi Kang, Philip Kauders, June Kim, Xiaoying Lin, Michael Luo, Max Mihm, Gregory Muenzen, Christopher Nieves, Collin Raymond, Alex Rees-Jones, John Schemitsch, Nathaniel Schorr, Dennis Shiraev, Nichole Szembrot, Russell Toth, Elizabeth Truax, and Ryan Yamada for their research assistance. We thank the National Institute on Aging (grants P30-AG012810 and P01-AG005842) for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Daniel J. Benjamin, James J. Choi, and Geoffrey W. Fisher. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Religious Identity and Economic Behavior
Daniel J. Benjamin, James J. Choi, and Geoffrey W. Fisher
NBER Working Paper No. 15925
April 2010, Revised January 2012
JEL No. D64,D71,D90,G11,H41,J22,Z12

ABSTRACT

We randomly vary religious identity salience in laboratory subjects to test how identity salience contributes to six hypothesized links from prior literature between religious identity and economic behavior. We find that religious identity salience makes Protestants increase contributions to public goods. Catholics decrease contributions to public goods, expect others to contribute less to public goods, and become less risk averse. Jews more strongly reciprocate as an employee in a bilateral labor market gift-exchange game. We find no evidence of religious identity salience effects on disutility of work effort, discount rates, or generosity in a dictator game.

Daniel J. Benjamin
National Bureau of Economic Research
1050 Massachusetts Avenue
Cambridge, MA 02138
and NBER
daniel.benjamin@gmail.com

Geoffrey W. Fisher
California Institute of Technology
geoffrey.w.fisher@gmail.com

James J. Choi
Yale School of Management
135 Prospect Street
P.O. Box 208200
New Haven, CT 06520-8200
and NBER
james.choi@yale.edu

Scholars have long hypothesized that religious identities affect economic outcomes. However, causal inference about the effect of religion has been hampered by the difficulty of identifying exogenous variation in religious identities. Religious affiliation—even when inherited from one’s parents—is likely to be correlated with many unobserved factors (e.g., social network composition, educational quality) that have causal impacts on economic behavior.

In this paper, we use a technique from experimental psychology to *create* exogenous variation in one channel through which religious identities could affect behavior: the salience of religious identity to the person affiliated with that identity. Religious identity will be salient at least some of the time to anybody who belongs to a religious category, so laboratory identity salience manipulations allow us to measure how behavior is likely to be affected by identity affiliation during those times.

We examine how religious identity salience might contribute to six hypothesized links between religious identity and economic behavior found in the previous literature:

- H1.* Contributions to public goods are affected positively by Protestantism and negatively by Catholicism (Putnam, 1993; La Porta et al., 1997).
- H2.* Trust is affected positively by Protestantism and negatively by Catholicism, and trust is the mechanism explaining the relationship between religion and public goods contributions (Putnam, 1993; La Porta et al., 1997).
- H3.* Financial risk-taking is affected by religion. Kumar, Page, and Spalt (2011) argue that risk-taking is reduced by Protestantism, and Hoffmann (2000) notes that the Catholic Church has often been criticized by other religious bodies for promoting gambling through bingo and other charitable gaming events. Hilary and Hui (2009) argue instead that risk-taking is reduced by both Protestantism and Catholicism.
- H4.* Thrift and capital accumulation are promoted by religion (Weber, 1930; Guiso, Sapienza, and Zingales, 2003; Barro and McCleary, 2003, 2006).
- H5.* Generosity is increased by religion (Friedrichs, 1960; Batson, Schoenrade, and Ventis, 1993).
- H6.* Work ethic is increased by religion, and especially by Protestantism (Weber, 1930; Barro and McCleary, 2003, 2006).

To operationalize these hypotheses in a laboratory setting, we map them into seven outcomes that can be measured using standard experimental methods. The first five outcomes are contributions to a laboratory public good, expectations about *others'* contributions to the laboratory public good, elicited risk aversion, elicited discount rates, and the amount given away in a dictator game. Since work ethic could be interpreted as a low disutility of effort or as a high reciprocity motive toward an employer when labor contracts are incomplete, we measure two outcomes corresponding to each of these concepts: the number of anagrams attempted at a constant piece-rate wage, and the rate at which effort expended as a worker rises with wage offers in a labor market bilateral gift-exchange game.

The results from the public goods game indicate that identity salience effects may contribute to the cross-sectional correlations between religion, contributions to public goods, and trust documented by Putnam (1993) and La Porta et al. (1997). We find that religious identity salience increases Protestants' contributions to the laboratory public good, while it causes Catholics to decrease contributions. Catholic identity salience decreases expectations of others' contributions to the public good, and it has no statistically significant effect on contributions once its impact on trust is controlled for. Protestant identity salience does not appear to affect trust, however, suggesting that its positive impact on public good contributions does not operate through a trust channel.

In accordance with the correlation between religion and financial risk-taking found by Kumar, Page, and Spalt (2011), we find that religious identity salience causes Protestants to become relatively more risk averse than Catholics. But in our data, this is because Catholic identity salience *increases* risk-taking, not because Protestant identity salience decreases it.

We find no evidence for Judeo-Christian identity salience effects on discount rates and generosity. Nor do we find Christian work ethic effects.¹ However, we do find a work ethic effect for Judaism that we had not hypothesized; among Jews, priming religion increases the rate at which workers increase their effort in response to higher wages in the gift-exchange game.

Why does identity salience affect behavior? A leading perspective in the psychology literature is "self-categorization theory," which has a long history (James, 1890; Turner, 1985). According to this theory, each person belongs to multiple social categories, such as religion,

¹ The lack of Christian work ethic effects and the divergence between Protestant and Catholic public good contribution effects may be consistent with the survey evidence of Arruñada (2010).

gender, and occupation, each of which has its own set of norms—prescriptions about how a member of a category *should* behave. Behavior in a given moment is more powerfully affected by the norms of categories that are salient than the norms of categories that are not salient. If an environmental cue, or a “prime,” makes a certain category temporarily more salient, behavior shifts towards the salient category’s norms. To clarify this self-categorization theory interpretation, we describe a simple economic model of identity salience effects. According to the model, priming effects reveal the sign of the marginal impact of category norms on steady-state preferences. The model also highlights a number of reasons why the magnitude of identity salience effects, as well as any failure to find an identity salience effect, are difficult to interpret.

Other researchers have manipulated religious identity salience and measured changes in subsequent generosity in a dictator game (Shariff and Norenzayan, 2007; Ahmed and Salas, 2011), risk choices (Hilary and Hui, 2009), willingness to take pamphlets from a charity (Pichon, Boccato, and Saroglou, 2007), cooperation in a prisoner’s dilemma game (Horton, Rand, and Zeckhauser, 2010; Ahmed and Salas, 2011), cheating on a laboratory task (Randolph-Seng and Nielsen, 2007; Mazar, Amir, and Ariely, 2008), and work effort on an unsolvable anagram task (Toburen and Meier, 2010). Our outcome measures are chosen specifically to test existing hypotheses in the economics literature. To our knowledge, this paper is the first to measure the impact of religious priming on public good contributions, expectations of others’ public good contributions, gift-exchange choices, and discount rates. When our outcome measures overlap with the prior literature, our results are distinguished by the fact that they are elicited under incentive compatibility and using much larger sample sizes. In addition, unlike previous research, we measure separate effects for each religion.

In related experimental economics work, Eckel and Grossman (2005), Hoff and Pandey (2006, 2011), Charness, Rigotti, and Rustichini (2007), Chen et al. (2011), and Chen and Chen (forthcoming) find that group identity effects on behavior are increasing in the salience of group membership.

Our paper proceeds as follows. Section I describes the pilot experiment we used to confirm that our priming instrument has the desired effect on religious identity salience. Section II describes the methodology of our main experiment. Section III presents the main experiment’s empirical results, including a joint hypothesis test to address concerns about Type I error due to multiple hypothesis testing. Section IV describes a theoretical framework for making inferences

about identity norms if self-categorization is the mechanism that generates our identity salience effects. Section V concludes.

I. Validating the Priming Instrument

The priming instrument, first used by Shariff and Norenzyan (2007) to study the effect of priming religious concepts, is a sentence-unscrambling task where subjects are asked to drop the irrelevant word in a five-word group and rearrange the remainder to form a four-word sentence. For example, “yesterday it finished track he” becomes “he finished it yesterday.” Each subject unscrambles ten sentences.

The sentences vary depending on whether the subject is in the religion-salient condition or the control condition. Five of the sentences unscrambled by religion-salient subjects contain religious content. These five sentences are: “she felt the spirit,” “the dessert was divine,” “give thanks to God,” “the book was sacred,” and “prophets reveal the future.” None of the control subjects’ sentences contain religious content.²

We recruited 91 students at the University of Michigan for a pilot experiment to confirm that the priming instrument increases the salience of religious identity. Subjects were randomly assigned to complete the religion-salient task or the control task. Subjects were not aware that this task differed across subjects. Immediately after the sentence unscramble, the questionnaire asked: “What five aspects of your identity (such as ‘male/female’ or ‘college student’) are most important to you?” Forty-seven percent of subjects in the religion-salient condition listed a religious identity in response, compared to only 25 percent of subjects in the control condition.³ This difference is significant at the 3 percent level (z -statistic = 2.16).

² An advantage of this priming instrument is that it is subtle; compared with blatant primes, subtle primes more reliably cause behavior to conform to norms (Wheeler and Petty, 2001), which aids in interpreting our results within the framework described in Section IV. Word-based puzzles are also used as primes in the economics literature by Matthey (2010), who finds that priming the concept of material achievements increases risk-taking, and by Drouvelis, Metcalfe, and Powdthavee (2010), who find that priming cooperative concepts increases public good game contributions. A potential disadvantage of this priming manipulation for our purposes is that “spirit” and “divine” may be less relevant to Jewish identity than they are to Catholic and Protestant identities, perhaps making our manipulation weaker for Jewish participants. Since eliciting subjects’ religion before the prime might have caused all subjects to become primed, we could not vary the manipulation by religion.

³ The fact that only 47 percent of subjects in the religion-salient condition listed a religious identity should not be interpreted as showing that the priming instrument is ineffective at making religion salient in half of all religious subjects. In our main experimental sample of Cornell students, about one-third identified themselves as atheists or agnostics. Although we did not elicit religious affiliations from our Michigan subjects (other than the ones who listed their religious affiliation as one of their top five identities), if a third of them were atheists or agnostics, then

We were concerned, however, that even if the prime succeeds in making religious identity more salient, it might *also* make related identity categories more salient. In that case, any priming effects we observe might be due to these other identity categories. To address this possibility, we tested whether the identity salience manipulation affected how frequently subjects mentioned any other identity besides religion. In descending order of frequency, subjects listed: gender, student, nationality/ethnicity, personality trait, family relationship, professional interest, non-family relationship, age group, activity group, political orientation, physical feature, and socioeconomic class. For each of these other identities, the p -value for a test of differences across conditions in the proportion of subjects who mentioned it varied from 25 to 97 percent. This evidence suggests that the priming instrument increases religious category salience without affecting the salience of other categories that are possibly correlated with religious categories.

II. Main Experiment Procedure

Participants in the main experiment were 827 Cornell University students. To avoid making religious identity salient to all the subjects, we did not mention in our recruiting materials that we were looking for people of particular religions. Sessions were administered by computer, most using the program z-Tree (Fischbacher, 2007) but some using VBA for Microsoft Excel, depending on which dependent variables we measured in the session.

Within each experimental session, we randomly assigned subjects to complete the religion-salient or control sentence unscramble. Subjects were not aware that this task differed across subjects. After completing the sentence unscramble, they participated in strategic games, incentive-compatible preference elicitations, and an anagram-solving task. We describe the games, elicitations, and anagram task below. Subjects were told at the beginning of the experimental session that any interactions they had with other subjects would be anonymous, one-shot interactions. In order to avoid excessively long sessions, each subject engaged in only a subset of the post-unscramble tasks.⁴ The order in which the post-unscramble tasks appeared varied across sessions.⁵

about 70 percent of primed religious subjects listed religion as one of their top five identities. Of course, the top-five measure does not detect instances where religious identity became more salient but did not ascend to the top five.

⁴ Two hundred fifteen subjects were in sessions that administered only the following four sections after the sentence unscramble: the dictator game, the public goods game, the risk preference elicitation, and the time preference elicitation. Three hundred forty-eight subjects were in sessions that administered only the following three sections after the sentence unscramble: the dictator game, the public goods game, and the labor market gift-exchange game.

A. Public goods game

We measured the willingness to contribute to a public good by assigning each subject to a group of four and endowing him or her with \$1.⁶ Subjects could contribute any fraction of their dollar to a group account, which is the laboratory public good. Contributions were doubled and then distributed evenly among the four group members. Subjects kept any money that they did not contribute. Total group earnings are maximized (at \$2 per group member) if each member contributes his or her entire dollar to the group account. However, in the absence of other-regarding preferences, it is a dominant strategy to contribute nothing, since the private return on a contribution is –50 percent.

Before eliciting subjects' own contributions, we asked subjects to give their best guess of how much the other three members of their group would contribute on average—a measure of trust.⁷ In order to keep the game's instructions relatively simple, we did not make subjects' payments depend upon the value of this guess.

Existing evidence suggests that behavior in laboratory public goods games like ours is associated with contributions to public goods outside the laboratory. Carpenter and Seki (forthcoming) find that cooperativeness in a laboratory public good game is positively correlated with the productivity of Japanese fishermen who pool their catch with other fishermen, and Fehr and Leibbrandt (2008) find that Brazilian fishermen who contribute more to a laboratory public good are less likely to over-exploit their common fishing ground. Laury and Taylor (2008) find

One hundred twenty-one subjects were in sessions that administered only the following three sections: the time preference elicitation, the risk preference elicitation, and the anagram work ethic task. One hundred forty-three subjects participated in sessions that administered only the anagram work ethic task. Sessions occurred between April 2008 and December 2009.

⁵ We have examined how the strength of the priming effect varied with the length of time elapsed since the priming manipulation. The effects on Jewish gift-exchange reciprocity and Catholic risk aversion may weaken over time, while the effects on Protestant and Catholic public goods contributions and Catholic trust appear to strengthen over time, but we put little weight on these possible trends because very few of the interactions between the priming effect and task order are statistically significant.

⁶ Marwell and Ames (1979) were the first to run this type of experiment to study the willingness to contribute to public goods.

⁷ We use the term “trust” in the sense of La Porta et al. (1997, p. 333): trust is when “people expect certain fair or cooperative behavior of their opponents even when they do not expect to see them again.” See also Gächter, Herrmann, and Thöni (2004) for an explication of how trust applies to a simultaneous-move public goods game.

that the amount contributed to a laboratory public good is positively correlated with the willingness to contribute money to a local tree-planting organization.⁸

B. Risk aversion elicitation

We elicited small-stakes risk preferences by asking participants to make six binary choices between \$1 for sure and a 50 percent chance at a larger amount, ranging from \$1.60 to \$3.60. The outcome of each gamble was independent, and we paid subjects for all six choices in accordance with their stated preferences; for example, if a subject selected the gamble if and only if winning the gamble yielded \$3.60, then she would have a 50 percent chance of earning \$5 and a 50 percent chance of earning $\$5 + \$3.60 = \$8.60$ in this section. We measured larger-stakes risk preferences with six analogous choices, where the monetary amounts were 100 times larger than the small-stakes risk choice amounts and there was only a small chance that the subject's choices would be implemented for payment.⁹ The section's instructions made it clear that the questions were not intended to evaluate subjects normatively: "It's important to keep in mind that there are no right or wrong answers here. Which choice you make is a matter of personal preference." We used this same wording again in the instructions for the discount rate elicitation.

Risk aversion measures derived from incentive-compatible experimental choices such as ours are highly correlated with measures from hypothetical choices, which in turn predict risky behaviors such as smoking, drinking, failing to hold insurance, holding stocks rather than Treasury bills, being self-employed, switching jobs, and moving residences.¹⁰

C. Discount rate elicitation

We measured time preferences by asking participants to make 12 binary choices between receiving \$10 now and receiving some larger amount one week from now, and another 12 binary

⁸ However, Laury and Taylor (2008) find that this relationship is driven by the most selfish subjects, and that differences in contributions among less selfish subjects do not reliably predict differences in contributions to the field public good. This may be because the choices about contributing to the laboratory public good and the field public good occurred close together in time. Virtuous behavior tends to encourage more selfish behavior subsequently due to the "licensing effect" (Khan and Dhar, 2006); because an individual has bolstered his self-concept as a virtuous person through the initial virtuous act, he is freed to act selfishly afterwards.

⁹ One of the six large-stakes risk choices was randomly chosen to be paid out if the subject could correctly predict two spins of a roulette wheel, which implied a one in 1,444 chance of one large-stakes risk choice being implemented.

¹⁰ See Barsky et al. (1997), Guiso and Paiella (2008), Dohmen et al. (2005), and Sahm (2007).

choices between receiving \$10 one week from now and receiving some larger amount two weeks from now. The larger delayed amounts ranged from \$10.10 to \$15. After a subject made these choices, we randomly selected one of the 24 choices and paid the subject according to his or her stated preference in this choice. All payments were made by a check given to the participant immediately following the experiment. Delayed payments were implemented via post-dated check.¹¹

Our approach to measuring time preferences is standard (Frederick, Loewenstein, and O'Donoghue, 2002). Similar measures predict variation in discounting-related behaviors such as drug addiction, cigarette smoking, excessive gambling, use of commitment savings devices, borrowing on installment accounts and credit cards, rapid exhaustion of food stamps, delayed application to an MBA program, and defaulting on loans.¹²

D. Dictator game

In our implementation of the dictator game (Kahneman et al., 1986; Forsythe et al., 1994), we endowed each subject with \$1 and randomly assigned him or her to another participant in the session. The subject could choose to give any portion of that \$1 to the other subject. A profit-maximizing individual would keep the entire dollar for himself, so the amount given away is a measure of pure altruism. Benz and Meier (2008) find that generosity in laboratory dictator games is positively correlated with charitable giving outside the laboratory.

E. Labor market tasks

Work ethic can be interpreted in terms of economic theory as relating to an individual's disutility of effort, which determines the willingness to exert a contracted-upon amount of effort at a given wage rate, or the strength of the reciprocity motive toward an employer that causes a worker to supply more effort in response to a higher wage when the labor contract is incomplete. We measure identity effects on both variables.

To measure identity effects on the disutility of effort, we asked subjects to solve as many four-letter anagrams as they could in a five-minute period. We paid participants five cents per

¹¹ If the subject received a delayed payment for this section, then earnings from other sections were paid through a separate check that was immediately cashable.

¹² See Fuchs (1982), Bickel, Odum, and Madden (1999), Kirby, Petry, and Bickel (1999), Petry and Casarella (1999), Kirby and Petry (2004), Shapiro (2005), Ashraf, Karlan, and Yin (2006), Meier and Sprenger (2010), and Reuben, Sapienza, and Zingales (2009).

correctly solved anagram. Because the piece-rate wage we offered is a contingent payment, subjects' reciprocity norm should not have been strongly activated in this task. Hence, the quantity of effort supplied by subjects equates the marginal cost of effort with the marginal benefit of the expected payment from exerting anagram-solving effort. Shifts in the amount of effort exerted across salience conditions reflect shifts in the marginal effort cost function.¹³ By restricting the task length to five minutes, we reduced the chances that primed subjects' religious identity salience would fade during the task.

We measured work provision in an incomplete contractual setting by running a labor market bilateral gift-exchange game (Fehr et al., 1998), modeled on one of the implementations of Charness, Frechette, and Kagel (2004). We paired subjects together and told them that one of them would play the role of the Manager and the other the role of the Employee. We described the roles in these terms so that subjects would be more likely to use norms applicable to labor markets in their choices. After playing once, subjects played the game again, but this time in the opposite role and with a different partner.

In the first stage of the game, the Manager paid a wage to the Employee between \$0 and \$4 that is a multiple of 50 cents. The Manager could not change the wage later. In the second stage, the Employee saw the wage and chose a work quantity to supply that is an integer between 1 and 10, inclusive. Employees were told that the Manager would be shown their work quantity choice. The Employee's cost of work provided was an increasing convex function of work quantity: \$0.00, \$0.04, \$0.08, \$0.16, \$0.24, \$0.32, \$0.40, \$0.48, \$0.60, and \$0.72 as work quantity rose from 1 to 10. The Employee's earnings were the wage received minus the cost of work provided. The Manager's earnings were $(\$4 - \text{wage paid}) \times \text{work quantity provided by the Employee} \div 10$.

To facilitate calculation, we provided subjects a lookup table that showed the Manager and Employee's earnings at each wage and work quantity combination. Managers simply chose one wage, while Employees indicated a contingent work supply choice for each possible wage. After observing the Managers' wage choice, we implemented the Employees' work supply

¹³ In mathematical notation, let e be the quantity of effort supplied, with units normalized so that the expected number of anagrams solved in the five minutes allotted is equal to e . Let p be the payment in the numeraire good per anagram solved, and let the increasing convex function $c(e)$ denote the utility cost of supplying effort. The subject's utility function is $U(e; p) = ep - c(e)$. As long as the optimal amount of effort e^* is interior, it is characterized by the first-order condition, $c'(e^*) = p$. Since p is the same across the treatment and control groups, changes in effort supply can be attributed to changes in the c' function.

choice based on that wage. The profit-maximizing strategy for the Employee is to always supply the minimum amount of work, since the Manager is unable to contract on effort and has no opportunity to punish the Employee for shirking. If the Manager believes the Employee is a profit maximizer, his or her own profit-maximizing response is to offer a \$0 wage. Despite these equilibrium predictions, the prior literature finds that, consistent with the presence of a reciprocity norm, Employees usually supply positive effort that is increasing in the wage offer. Managers anticipate this norm and rarely offer the minimum possible wage.

Because the Employee's work quantity is the choice of a number rather than effort in a real work task, and because the cost of providing this work quantity is determined by a function that is the same for all subjects, differences in work provided across salience conditions are driven solely by changes in the strength of subjects' reciprocity norm and not by changes in the utility cost of effort.

F. Post-experimental questionnaire

At the end of the session, after payoffs had been revealed, subjects completed a questionnaire that collected information about their demographic characteristics, beliefs about the experiment, and religious beliefs. We also included numerous decoy questions to mask the purpose of the study, so that subjects would not contaminate future subjects by telling them that we were running an experiment about religion. We discuss responses to the relevant questions in further detail in Section III.

III. Main Experiment Results

A. Sample selection

Our sample consists of 264 Protestant/other Christians (whom we refer to collectively as "Protestants"), 199 Catholics, 95 Jews, and 269 atheists/agnostics.¹⁴ We are not aware of existing hypotheses in the literature regarding how atheism or agnosticism might affect economic behavior. Moreover, it is possible that, among subjects who were raised in a faith and have subsequently abandoned it, religious identity salience makes salient the religious affiliation in

¹⁴ Since we analyze results separately by religion, it would be problematic if the priming manipulation affected how participants categorized themselves, e.g., causing marginal atheist/agnostics to identify themselves with the religion in which they were raised. We find that religious affiliations are almost perfectly balanced across treatment and control groups. In the treatment group, there are 130 Protestants, 101 Catholics, 43 Jews, and 136 atheists/agnostics, compared with 134, 98, 52, and 133 in the control group, respectively.

which the subject was raised. For these reasons, we do not attempt to interpret religious salience effects on atheists and agnostics. Nonetheless, we report results from all of the data we collected, including for this group.

In the post-experimental questionnaire, we asked subjects, “What do you think this study is about?” To avoid estimating treatment effects that are driven by experimenter demand effects, the above sample of 827 excludes four subjects who guessed that the study had something to do with religion.¹⁵ In addition, we have dropped four subjects who left more than half of the sentence-unscramble responses blank, since they were likely not to have been properly primed. Our results are not sensitive to including these dropped subjects.

B. Public goods game

Panel A of Table 1 shows coefficients from regressing the amount contributed to the public good on a dummy for being in the religion-salient condition. The constant coefficients indicate that among unprimed subjects, Catholics contribute the most on average, agnostics and atheists contribute the least, and Protestants and Jews are in the middle. However, selection into our sample of Cornell students is not random. And even if our sample were representative of each religion’s members, the many unobserved variables that are correlated with religious affiliation would prevent us from inferring any causal effects of religion by simply comparing subject choices across religions. We instead turn to comparisons between the treatment and control groups within each religion, which identify the causal impact of religious identity salience using randomized variation.

We find that Protestant identity salience increases the willingness to contribute in a public goods game while Catholic identity salience suppresses it. Protestants for whom religious identity is salient contribute 15 cents more to the public good than control Protestants, whereas primed Catholics decrease their contributions by 18 cents. Jewish, atheist, and agnostic subjects’ contributions are not significantly affected by the prime.

¹⁵ In the post-experimental questionnaire, we also asked subjects whether they believed that their experimental choices would affect their payments as the instructions specified. Ninety percent of subjects reported believing these payment promises. Among the subjects who participated in the dictator, public goods, and gift exchange games, 88 percent reported believing that their choices would affect *other* participants’ payments exactly as we had specified. Most economics experiments do not measure payment promise credibility. Our subjects’ rates of belief are much higher than those found in the Benjamin, Choi, and Strickland (2010) sample.

Putnam (1993), La Porta et al. (1997), and Ruffle and Sosis (2007) argue that the channel through which religion affects public goods provision is trust. Among Catholics, our data support the hypothesis that identity salience affects public goods contributions by changing the propensity to trust. Panel B of Table 1, which contains regressions of subjects' expectations of other group members' average contribution on a religion-salient dummy, indicates that relative to control Catholics, primed Catholics expect the average member of their group to contribute 12 cents less. However, there is no significant effect of priming on Protestant, Jewish, or atheist/agnostic expectations. This suggests that Protestant identity salience's positive effect on public good contributions in Panel A is not driven by a change in Protestants' trust.

Panel C shows that subjects' reported expectations, despite being unincentivized, are nonetheless strongly predictive of their behavior. The coefficients in this panel are from a regression of contributions on a religion-salient dummy and the subject's expectation of other group members' average contribution. We find that subjects' own contributions increase almost one-for-one with their expectations.¹⁶ Once expectations are controlled for, primed Catholics do not contribute significantly less than unprimed Catholics, suggesting that the Catholic priming effect on contributions operates largely through Catholicism's negative effect on trust. In contrast, primed Protestants contribute a highly significant 12 cents more than unprimed Protestants, even after controlling for expectations. The Protestant priming coefficient is statistically distinguishable from the Catholic priming coefficient at the 5 percent level.

C. Risk aversion

In the risk preference regression, our dependent variable is the minimum risk premium—that is, the expected return offered by the gamble in excess of the risk-free return—that the subject requires to accept the gamble. For example, if $X = \$2.40$ is the smallest X for which a subject would choose to gamble for a 50 percent chance of receiving X rather than accept a sure \$1, then the reservation risk premium is $(2.40 \times 0.5 - 1)/1 = 0.20$.¹⁷ Because we observe binary

¹⁶ In untabulated regressions, we find that this relationship does not differ significantly between primed and unprimed subjects in any of the religious groups.

¹⁷ In this formula, we treat the risk choice as investing a \$1 endowment into either the risk-free or risky asset. Our decision to use the “minimum” risk premium as the dependent variable instead of the maximum risk premium that induces a safe choice only matters when participants' risk choices are inconsistent, switching between choosing the safe option and the risky option at more than one value of X . Since 95 percent of participants behaved consistently, our results are virtually identical if we use the “maximum” risk premium instead. Similarly, our choice to use the

choices over only a finite number of risk premia, we use an interval regression (Stewart, 1983), which is a generalization of the tobit estimator that accommodates dependent variable values that are not precisely observed but are known to lie somewhere within an interval. We observe two risk premia for each subject—one for the small-stakes gamble and one for the larger-stakes gamble—so each subject appears in the regression twice, and we cluster standard errors by subject (Froot, 1989; Rogers, 1993).

Table 2 shows that priming religious identity causes the average risk premium required to entice agnostics and atheists to forego a sure payout to fall significantly by 12 percentage points. For the small-stakes gamble, this represents a fall from 16 percent to 4 percent. This effect is driven almost entirely by the agnostics. We also find evidence that religious identity salience causes Catholics' risk premium to fall by 11 percentage points. This is a drop from 21 percent to 10 percent for the small-stakes gamble. The effect is significant only at the 10 percent level in this regression specification, but if we additionally control for a gender dummy (which is highly significant), the priming effect point estimate becomes significant at the 5 percent level. We find no significant identity salience effects on risk-taking for Protestants and Jews.

Our results are consistent with identity salience contributing to Kumar, Page, and Spalt's (2011) conclusion that Catholics are more risk-seeking than Protestants. But whereas they argue that the difference is due to Protestantism depressing risk tolerance, in our data, the wedge identity salience creates is instead driven by Catholics becoming more risk-seeking.¹⁸

D. Discount rate

In the time preference regression, our dependent variable is the log of the minimum continuously compounded weekly interest rate that the subject requires to choose the later payment. That is, we apply the log operator once to transform the reservation gross interest rate into the continuously compounded net interest rate, and then we apply the log operator again. The second application of the log operator causes the estimator we describe below to assume that reservation continuously compounded interest rates are conditionally log-normal, thus ruling out

“minimum” interest rate in our time preference regressions in Section IV.D does not matter because 97 percent of participants had a unique interest rate switch point.

¹⁸ Although it is not the focus of their paper, Hilary and Hui (2009) find suggestive evidence that priming religion increases risk aversion in a sample of 120 undergraduates. It is difficult to directly compare their results with ours because they report only full-sample results and not the religious composition of their sample, they report only *p*-values (0.20 and 0.08 for their two regressions) rather than effect sizes, they use a different priming instrument, and their subjects made only two risk choices that are both hypothetical.

negative discount rates. Each subject appears in the regression twice because we have two discount rate observations: now versus one week in the future, and one week versus two weeks in the future. Therefore, we cluster our standard errors by subject.

As in the risk preference elicitation, we only observe binary choices at a finite number of interest rates. Therefore, we use interval regressions to estimate subjects' reservation interest rates. In the interest rate regressions that follow, if the coefficients imply that a certain set of explanatory variable values is associated with a mean log continuously compounded interest rate of $\hat{\mu}$, then the median continuously compounded interest rate is $\exp(\hat{\mu})$.

Table 3 presents the regression evidence on how priming religious identity affects discount rates. The explanatory variables are a dummy for being in the religion-salient condition and a dummy for the intertemporal choice being between payments deferred for one week versus two weeks. In no case do we find that religious identity salience affects discount rates in a statistically significant way. Moreover, the point estimates of the priming effect are positive for all religions. If religion promotes thrift and capital accumulation (Weber, 1930; Guiso, Sapienza, and Zingales, 2003; Barro and McCleary, 2003 and 2006), it appears to do so through channels other than identity salience.

E. Dictator game

Despite having incentives to keep their entire \$1 endowment for themselves, dictators in our dictator game usually gave away a positive amount of money to the subject with whom they were paired, although the proportion given away was far less than half on average. This is a typical result for dictator game experiments (e.g., Forsythe et al., 1994). Table 4 contains coefficients from regressing the amount of money given away on a dummy for being in the religion-salient condition. The constant terms indicate that when unprimed, the average amount given away is between 14 and 22 cents, depending on the religious group.

The coefficients on the religion-salient dummy show that the prime does not induce any of the religious groups to give away significantly more money. The Protestant and Catholic priming effect point estimates are in fact negative.

Thus, we fail to replicate the Shariff and Norenzyan (2007) and Ahmed and Salas (2011) findings that priming religion increases generosity in a dictator game, even though we use the

same priming instrument as Shariff and Norenzayan and a similar one as Ahmed and Salas.¹⁹ Neither paper reports results separately by religion, but a rough meta-analysis using full-sample results points to a very small, positive, and insignificant effect overall. Assuming independence between Shariff and Norenzayan's Study 1 ($\beta = 0.238$, $s.e. = 0.0645$, $N = 50$), their Study 2 ($\beta = 0.200$, $s.e. = 0.0810$, $N = 50$), Ahmed and Salas' experiment ($\beta = 0.0714$, $s.e. = 0.0217$, $N = 224$), and our experiment ($\beta = -0.034$, $s.e. = 0.0214$, $N = 563$), the combined estimate of the priming effect on dictator game giving for a \$1 endowment is \$0.021 with a standard error of \$0.016 ($p = 0.189$).²⁰

F. Labor market tasks

In the anagram solving task, our subjects on average supplied answers for around 30 anagrams during the five minutes allotted to them. But contrary to the hypothesis that religious identity salience will decrease the cost of work effort, the regression in Table 5 of anagrams attempted on a dummy for being in the religion-salient condition shows that priming religious identity does not significantly affect the number of anagrams attempted among any of our religious groups. The point estimate of the priming effect on Protestants, who were the object of Weber's (1930) Protestant work ethic hypothesis, is in fact negative, although this is not significantly different from zero. In untabulated analysis, we have instead used the number of anagrams correctly answered as the dependent regression variable and find qualitatively similar results.

Labor supply effects in the bilateral gift-exchange game are summarized in Figure 1, which plots the average amount Employees chose to expend on work-related costs for each wage offer. Among all religious groups and experimental conditions, minimal labor is supplied at a \$0

¹⁹ Our results may differ because our subjects performed several tasks after the priming manipulation, while Shariff and Norenzayan's subjects played only the dictator game, and Ahmed and Salas's subjects played the dictator game immediately after the prime. However, when we restrict our analysis to the sessions where the dictator game was the very first task after the prime, we still do not find an effect in any religious group. Pooling across religious groups to make our results maximally comparable to Shariff and Norenzayan's, we find a coefficient of -0.016 and a standard error of 0.035 in this subsample of 222 individuals.

²⁰ To conduct this calculation, we divide Shariff and Norenzayan's and Ahmed and Salas's coefficients and standard errors by their stake sizes in order to put all of the findings in units of allocating a \$1 endowment. In analyzing Shariff and Norenzayan's Study 2, we ignore their second treatment group (which received a non-religious prime). Of course, it is possible that there are unpublished data that would be relevant for this meta-analysis, but we do not have access to those findings.

wage offer.²¹ The amount Employees are willing to expend on work-related costs rises nearly linearly with wage offers between \$0 and \$3.50, indicating a strong reciprocity norm. In measuring reciprocity, we will ignore labor supply at a \$4 wage because at this wage, the Manager's earnings are always \$0 regardless of how much labor the Employee supplies. Therefore, any positive work-related expenditures by the Employee at a \$4 wage represents either intentional money burning as a gesture of goodwill and appreciation or confusion about the payoff formulas. Labor supply is usually lower at a \$4 wage than at a \$3.50 wage.

It is readily apparent from Figure 1 that the strength of reciprocity, as reflected in the slope of work expenditures with respect to wages between \$0 and \$3.50, increases greatly among primed Jews relative to unprimed Jews. Other religious groups' reciprocity does not appear to be nearly as affected by priming. We formally analyze the priming effect on Employees' reciprocity in Panel A of Table 6. The dependent variable in the regression is the subject-specific slope coefficient from a regression of the subject's work-related costs on wage offers from \$0 to \$3.50. Jewish identity significantly increases this slope from 0.05 to 0.11, but there are no significant effects for any other group.

Panel B of Table 6 examines the effect of priming religious identity on wages offered by Managers. There is to our knowledge no clearly articulated hypothesis in the prior literature about religion's effect on managerial wage offers, but we report these results for completeness.²² The regression of managerial wage offers on a dummy for being in the religion-salient condition shows no significant effects of religious identity.

G. Multiple hypothesis testing and Type I error

Although all of the main hypotheses we have tested correspond to existing hypotheses in the literature, the number of tests we have run is large, and many of the priming effects are insignificant. This raises the concern that our significant priming effects are chance artifacts arising from the large number of tested hypotheses.

²¹ The absence of positive religious identity salience effects on purely altruistic labor provision in the gift-exchange game at a \$0 wage is consistent with the lack of religious identity salience effects on pure altruism in the dictator game.

²² However, since one motivation for a Manager to pay a positive wage is that he or she expects the Employee to reciprocate the kindness, one might expect that hypotheses about religious identity effects on trust in the public goods context would also apply to Managerial wage offers. Our results are not consistent with this extension and instead suggest that Catholic identity salience effects on trust are specific to public good contexts.

To address this possibility, we test whether the priming effects we estimated on public good contributions, expectations of others' public good contributions, risk aversion, discount rates, dictator game generosity, anagrams attempted, and gift-exchange reciprocity are jointly equal to zero. We pool all of our observations into a single interval regression and control for outcome type dummies (e.g., a dummy for the dependent variable being a public good contribution amount, a dummy for the dependent variable being an amount given away in the dictator game, etc.), outcome type dummies \times religion dummies, and outcome type dummies \times religion dummies \times additional outcome-specific control variables. Interval regressions can accommodate dependent variables that are point observations as a special case. We allow the residual's variance to vary by religion \times outcome type and cluster standard errors by subject. This procedure causes every individual priming coefficient value we previously estimated to appear in a single coefficient vector, allowing us to run a Wald test for their joint equality with zero. We reject this hypothesis at $p = 0.006$.²³

H. Treatment interactions with belief in divine punishment and religious service attendance

In this subsection, we examine whether the priming effects we have identified as significant for Christians or Jews differ for subjects with a stronger belief in divine punishment or more regular religious service attendance. Barro and McCleary (2003, 2006) find that the positive association between religion and GDP growth appears to operate most strongly through belief in the existence of hell, which could be a powerful motivator of behavior. In contrast, they find a negative association between GDP growth and the frequency of religious service attendance. Although Barro and McCleary interpret this negative relationship as arising from religious service attendance being a proxy for real resources being diverted to religious activities rather than economic production, it is possible that identity effects operate differently for frequent attendees than infrequent attendees, even holding beliefs about the afterlife fixed.

Our post-experimental questionnaire asked a subset of our subjects to rate on a six-point Likert scale their agreement with the statement, "God punishes people for their sins." We normalize this variable so that within each religious group, it has a zero mean and unit variance. We also asked all subjects how often they attend religious services. Possible answers were

²³ We exclude the priming effect on Managers' wage offers from this joint hypothesis test because there is no strong *ex ante* hypothesis from the literature about how religion would affect this dependent variable. If we include it, the *p*-value for the joint hypothesis test is 0.015.

“never,” “less than once a month,” “once a month,” “a few times a month,” “once a week,” “a few times a week,” “once a day,” and “more than once a day.” We create an indicator variable for whether the subject’s attendance frequency is above the median for his or her religious group. Median attendance frequency is once a month for Protestants, less than once a month for Catholics and Jews, and never for atheists and agnostics.²⁴

Table 7 shows regressions where the explanatory variables include an interaction of the religion-salient dummy with either the strength of belief in divine punishment or with the indicator for greater-than-median frequency of religious service attendance. To keep the number of interaction regressions manageable, we limit attention to the dependent variables and religious groups where we found statistically significant main effects of priming. For the sake of brevity, we omit from the table regressions with Catholic trust as the dependent variable, since these results are similar to those from regressions with Catholic public good contributions as the dependent variable. We also omit regressions involving agnostics and atheists because interactions with religious belief or religious service attendance are difficult to interpret for this group.

We find no statistically significant interactions of religious identity effects with attendance frequency or belief in divine punishment among Protestants and Catholics. Among Jews, we find that the increase in Employee reciprocity in response to the religious prime occurs primarily among those who attend religious services relatively *infrequently*.

IV. A Theoretical Framework for Interpreting Priming Effects

To provide a perspective on why identity salience affects behavior and what can be inferred from it, we discuss a theoretical framework that formalizes the “self-categorization theory” perspective from the psychology literature (James, 1890; Turner, 1985). This framework, which is inspired by Akerlof and Kranton (2000), is developed in Benjamin, Choi, and Strickland (2010). Here we elaborate on it, especially emphasizing what *cannot* be inferred from identity salience effects, even if the self-categorization theory interpretation is taken as given.

²⁴ We find no systematic differences in either the belief or attendance variables across treatment and control groups, with the exception that primed Jews report substantially higher belief in divine punishment ($p < .05$). The coefficients on belief in divine punishment and its interaction with the treatment dummy are essentially zero in the gift-exchange reciprocity regression among Jews in Table 7, so differences in this belief across the Jewish experimental groups are unlikely to be driving the Jewish reciprocity treatment effect.

The basic idea is that priming a social category temporarily increases the strength of affiliation with that category. Stronger affiliation with a category causes behavior to shift towards that category's norms, so comparing primed and unprimed behavior allows us to infer something about what the category's norms are and how they affect steady-state behavior.

Let x be some choice, such as how much to contribute to a public good or how much to trust a stranger. An individual belongs to a social category C , such as Protestant, with strength $s \geq 0$. Let x_0 denote the individual's preferred choice in the absence of identity considerations, and let x_C denote the choice prescribed for members of social category C . The individual chooses x to maximize

$$U = -(1 - w(s))(x - x_0)^2 - w(s)(x - x_C)^2, \quad (1)$$

where $0 \leq w(s) \leq 1$ is the weight placed on the norm for social category C in the person's decision. We assume that $w(0) = 0$ and $w' > 0$. Deviating from the norm prescribed for one's category causes disutility that is increasing in s . We assume that s has a steady-state value \bar{s} but can be temporarily increased to $\bar{s} + \varepsilon$, where $\varepsilon > 0$, by a category prime that makes the category's norms more cognitively accessible and hence more influential on behavior.

The first-order condition of (1) gives the optimal action, $x^*(s) = (1 - w(s))x_0 + w(s)x_C$, which is a weighted average of the preferred action without identity considerations and the category norm. This condition yields several implications:

1. The higher the steady-state strength \bar{s} of the category affiliation, the closer x^* is to x_C in steady state.
2. A category prime also causes x^* to move closer to x_C . Thus, the behavioral effect of priming social category C reveals the marginal behavioral effect of increasing the steady-state strength \bar{s} of C . This makes priming manipulations a useful experimental procedure for studying how identity affects steady-state choices.
3. The sign of the priming treatment effect, $x^*(s + \varepsilon) - x^*(s) \approx (dx^*/ds)\varepsilon = w'(s)(x_C - x_0)\varepsilon$, depends on the sign of $x_C - x_0$. Even if the \bar{s} , x_0 , and $w(\cdot)$ of an experimental sample differ from those of the general population affiliated with C , the directional effects of

priming the sample will generalize to the population as long as $x_C - x_0$ has the same sign for both groups.²⁵

Psychologists have tested the category salience mechanism by priming social categories with norms that are believed to be known and confirming that choices shift towards these norms. For example, LeBoeuf, Shafir, and Belyavsky (2010) find that undergraduate subjects are more likely to prefer highbrow periodicals and films when their “scholar” social category is primed than when their “socialite” social category is primed. Chinese-American subjects make more collectivist choices when their Chinese social category is primed rather than their American social category. If subjects make a choice while one social category is salient, they are subsequently more dissatisfied with that choice if a different social category is salient at the time of post-choice evaluation than if the same social category remains salient. Reicher and Levine (1994) find that undergraduate science majors express more favorable attitudes towards practices such as animal vivisection when their scientist identity category is made salient, and Forehand, Deshpandé, and Reed (2001) find that making subjects’ ethnicity salient causes them to evaluate a same-ethnicity spokesperson more favorably. Inverting the approach in these papers by taking the validity of the category salience mechanism as given and priming categories with *unknown* norms, we obtain information about the norms from the resulting shift in choices.

Importantly, this framework also provides a way to think about the limitations of priming techniques for identifying norms:

4. Although the direction of the priming effect reliably identifies the sign of $x_C - x_0$, differences in the priming effect’s magnitude across people may arise through a number of channels. Assume without loss of generality that $x_C > x_0$. Priming will have a larger effect if the identity norm is more extreme (i.e., x_C is larger) or the person’s preferred action in the absence of identity considerations is more extreme in the opposite direction (i.e., x_0 is smaller). Priming will also have a larger effect if the salience manipulation is more effective at increasing identity salience for that particular person (i.e., ε is larger) or the person’s choices are more sensitive to a given change in identity salience (i.e., w' is

²⁵ The sign of $x_C - x_0$ might differ also because x_C in a sample differs from that in the general population affiliated with C . This could occur, for example, if identity norms differ across subgroups within a religious category—such as different denominations within Protestantism—and the mix of subgroups differs between the sample and the general population affiliated with that category.

larger). This latter difference can arise either because the w function has a different shape or because the person has a different steady-state \bar{s} , so that the points at which she evaluates the w function differ. Because it is difficult to interpret differences in priming effect magnitudes, we have focused on estimating the priming effect within a religious identity category, rather than a difference-in-difference analysis of priming effects across religious categories.

5. Priming will not reveal social category effects that operate exclusively through x_0 rather than x_C . For example, many Koreans enjoy eating spicy foods, and this preference would probably be weaker if they were not born into the Korean social category. However, if eating spicy food is simply a taste that Koreans develop because Korean culture exposes them to spicy foods as children (i.e., a shift in x_0 caused by category membership), and refraining from spicy foods is not thought to make one a “bad” or “disloyal” Korean (i.e., there is no x_C associated with the spiciness of a Korean’s diet, so the second additive term in equation (1) is missing), then priming the Korean social category will not cause Koreans to eat spicier foods.
6. There will be no difference between primed and unprimed choices in domains where the choice situation itself functions as a strong social category prime. For example, a Jew who is strictly kosher will have his Jewish identity primed whenever he is presented with the choice to eat pork. He will always refuse to eat the pork whether or not an additional Jewish category cue is present. Our intuition is that choice situations to which a person considers the social category norm to be strictly and unambiguously applicable are more likely to be primes in themselves.²⁶

Points 5 and 6 imply that null priming effects need not be interpreted as definitive evidence against the existence of social category effects in a particular domain.

V. Conclusion

In our data, we find that religious identity salience among Protestants causes increased contributions to public goods, while religious identity salience among Catholics causes decreased

²⁶ Thus, we hypothesize that priming the Jewish category *would* make a Jew who occasionally keeps kosher outside the Jewish holidays more likely to refuse pork, since she both places some weight on the kosher norm in her utility function and is not likely to be thinking about it during an eating situation without the prime.

contributions to public goods, lower expectations of others' contributions to public goods, and decreased risk aversion. Religious identity salience among Jews increases labor market reciprocity. However, we find no evidence that religious identity salience affects discount rates or purely altruistic generosity.

We conclude with two notes of caution in interpreting our findings. First, when we fail to find an effect of identity salience, it does not necessarily rule out the effect's existence. It is possible that we did not have sufficient statistical power to detect the effect, or that a stronger salience manipulation would have produced a detectable effect. It is also possible that the strength of our identity salience manipulation varies by religious group, making us more likely to find null effects in some groups than others. Hence, our positive findings merit greater emphasis than our null findings.

Second, while a key advantage of identity salience manipulations is that they can be randomly assigned, exogenously manipulating identity salience is not the same thing as exogenously manipulating religious affiliation. For example, making religious identity salient is a different experiment than raising a child in a different religious tradition. According to self-categorization theory, identity salience effects identify the marginal directional effect of religious identity norms. Thus, identity salience manipulations provide information regarding one channel—identity norms—through which religion might affect behavior. But regardless of the mechanisms responsible for our identity salience effects, religious identity will be salient outside the laboratory some of the time to members of a particular religious category, so identity salience manipulations allow us to measure how behavior is likely to be affected by identity affiliation during those times. Therefore, we believe that identity salience manipulations—interpreted appropriately and cautiously—can be a useful tool for learning about some of the effects of identity.

References

- Ahmed, Ali M., and Osvaldo Salas. 2011. "Implicit influences of Christian religious representations on dictator and prisoner's dilemma game decisions." *Journal of Socio-Economics* 40: 242-246.
- Akerlof, George A., and Rachel E. Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics* 115(3): 715-733.

- Arruñada, Benito. 2010. "Protestants and Catholics: Similar Work Ethic, Different Social Ethic." *Economic Journal* 120(547): 890-918.
- Ashraf, Nava, Dean S. Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Project in the Philippines." *Quarterly Journal of Economics* 121(2): 635-672.
- Barro, Robert J., and Rachel M. McCleary. 2003. "Religion and Economic Growth." *American Sociological Review* 68: 760-781.
- Barro, Robert J., and Rachel M. McCleary. 2006. "Religion and Economy." *Journal of Economic Perspectives* 20: 49-72.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* 112(2): 537-579.
- Batson, C. Daniel, Patricia A. Schoenrade, and W. Larry Ventis. 1993. *Religion and the Individual: A Social-Psychological Perspective*. New York: Oxford University Press.
- Benjamin, Daniel J., James J. Choi, and A. Joshua Strickland. 2010. "Social Identity and Preferences." *American Economic Review* 100(4): 1913-1928.
- Benz, Matthias, and Stephan Meier. 2008. "Do People Behave in Experiments as in the Field?—Evidence from Donations." *Experimental Economics* 11: 268-281.
- Bickel, Warren K., Amy L. Odum, and Gregory J. Madden. 1999. "Impulsivity and Cigarette Smoking: Delay Discounting in Current, Never, and Ex-Smokers." *Psychopharmacology* 146(4): 447-454.
- Carpenter, Jeffrey, and Erika Seki. Forthcoming. "Do Social Preferences Increase Productivity? Field Experimental Evidence from Fishermen in Toyoma Bay." *Economic Inquiry*.
- Charness, Gary, Guillaume R. Frechette, and John H. Kagel, 2004. "How Robust is Laboratory Gift Exchange?" *Experimental Economics* 7: 189-205.
- Charness, Gary, Luca Rigotti, and Aldo Rustichini. 2007. "Individual Behavior and Group Membership." *American Economic Review* 97: 1340-1352.
- Chen, Roy, and Yan Chen. Forthcoming. "The Potential of Social Identity for Equilibrium Selection." *American Economic Review*.
- Chen, Yan, Sherry Xin Li, Tracy Xiao Liu, and Margaret Shih. 2011. "Which Hat to Wear? Impact of Natural Identities on Coordination and Cooperation." University of Michigan working paper.

- Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner. 2005. "Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey." IZA Discussion Paper No. 1730.
- Drouvelis, Michalis, Robert Metcalfe, and Nattavudh Powdthavee. 2010. "Priming Cooperation in Social Dilemma Games." IZA Discussion Paper No. 4963.
- Eckel, Catherine C., and Philip J. Grossman. 2005. "Managing Diversity by Creating Team Identity." *Journal of Economic Behavior & Organization* 58: 371-392.
- Fehr, Ernst, and Andreas Leibbrandt. 2008. "Cooperativeness and Impatience in the Tragedy of the Commons." IZA Discussion Paper No. 3625.
- Fehr, Ernst, Erich Kirchler, Andreas Weichbold, and Simon Gächter. 1998. "When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets." *Journal of Labor Economics* 16(2): 324-351.
- Fischbacher, Urs. 2007. "z-tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10: 171-178.
- Forehand, Mark R., Rohit Deshpandé, and Americus Reed II. 2002. "Identity Salience and the Influence of Differential Activation of the Social Self-Schema on Advertising Response." *Journal of Applied Psychology* 87(6): 1086-1099.
- Forsythe, Robert, Joel L. Horowitz, N. E. Savin, and Martin Sefton, 1994. "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior* 6: 347-369.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40(2): 351-401.
- Friedrichs, Robert W. 1960. "Alter versus Ego: An Exploratory Assessment of Altruism." *American Sociological Review* 25(4): 496-508.
- Froot, Kenneth A. 1989. "Consistent Covariance Matrix Estimation with Cross-Sectional Dependence and Heteroskedasticity in Financial Data." *Journal of Financial and Quantitative Analysis* 24(3): 333-355.
- Fuchs, Victor. 1982. "Time Preference and Health: An Exploratory Study." In Victor R. Fuchs, editor, *Economic Aspects of Health*: 93-120. Chicago: University of Chicago Press.
- Gächter, Simon, Benedikt Herrmann, and Christian Thöni. 2004. "Trust, Voluntary Cooperation, and Socio-Economic Background: Survey and Experimental Evidence." *Journal of Economic Behavior & Organization* 55(4): 505-531.
- Guiso, Luigi, and Monica Paiella. 2008. "Risk Aversion, Wealth, and Background Risk." *Journal of the European Economic Association* 6(6): 1109-1150.

- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2003. "People's Opium? Religion and Economic Attitudes." *Journal of Monetary Economics* 50: 225-282.
- Hilary, Gilles, and Kai Wai Hui. 2009. "Does Religion Matter in Corporate Decision Making in America?" *Journal of Financial Economics* 93: 455-473.
- Hoff, Karla, and Priyanka Pandey. 2006. "Discrimination, Social Identity, and Durable Inequalities." *American Economic Review* 96(2): 206-211.
- Hoff, Karla, and Priyanka Pandey. 2011. "Making up people: Experimental evidence on identity and development from caste India." World Bank mimeo.
- Hoffmann, John P. 2000. "Religion and Problem Gambling in the U.S." *Review of Religious Research* 41: 488-509.
- Horton, John J., David G. Rand, and Richard J. Zeckhauser. 2010. "The Online Laboratory: Conducting Experiments in a Real Labor Market." NBER Working Paper 15961.
- James, William. 1890. *The Principles of Psychology*. New York: Henry Holt.
- Kahneman, Daniel, Jack Knetsch, and Richard H. Thaler. 1986. "Fairness and the Assumptions of Economics." *Journal of Business* 59: S285-300.
- Khan, Uzma, and Ravi Dhar. 2006. "Licensing Effect in Consumer Choice." *Journal of Marketing Research* 43: 259-266.
- Kirby, Kris N., and Nancy M. Petry. 2004. "Heroin and Cocaine Abusers Have Higher Discount Rates for Delayed Rewards than Alcoholics or Non-Drug-Using Controls." *Addiction* 99(4): 461-471.
- Kirby, Kris N., Nancy M. Petry, and Warren K. Bickel. 1999. "Heroin Addicts Have Higher Discount Rates for Delayed Rewards than Non-Drug-Using Controls." *Journal of Experimental Psychology* 128(1): 78-87.
- Kumar, Alok, Jeremy Page, and Oliver Spalt. 2011. "Religious Beliefs, Gambling Attitudes, and Financial Market Outcomes." *Journal of Financial Economics* 102(3): 671-708.
- La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert W. Vishny. 1997. "Trust in Large Organizations." *American Economic Review* 87: 333-338.
- Laury, Susan K., and Laura O. Taylor. 2008. "Altruism Spillovers: Are Behaviors in Context-Free Experiments Predictive of Altruism Toward a Naturally Occurring Public Good?" *Journal of Economic Behavior and Organization* 65: 9-29.
- LeBoeuf, Robyn A., Eldar Shafir, and Julia Belyavsky Bayuk. 2010. "The Conflicting Choices of Alternating Selves." *Organizational Behavior and Human Decision Processes* 111(1): 48-61.

- Marwell, Gerald, and Ruth E. Ames. 1979. "Experiments on the Provision of Public Goods. I. Resources, Interest, Group Size, and the Free-Rider Problem." *American Journal of Sociology* 84(6): 1335-1360.
- Matthey, Astrid. 2010. "The Influence of Primes on Reference States." *Games* 1(1): 34-52.
- Mazar, Nina, On Amir, and Dan Ariely. 2008. "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance." *Journal of Marketing Research* 45(6): 633-644.
- Meier, Stephan, and Charles Sprenger. 2010. "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics* 2(1): 193-210.
- Petry, Nancy M., and Thomas Casarella. 1999. "Excessive Discounting of Delayed Rewards in Substance Abusers with Gambling Problems." *Drug and Alcohol Dependence* 56(1-2): 25-32.
- Pichon, Isabelle, Giulio Boccato, and Vassilis Saroglou. 2007. "Nonconscious Influences of Religion on Prosociality: A Priming Study." *European Journal of Social Psychology* 37: 1032-1045.
- Putnam, Robert. 1993. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton, NJ: Princeton University Press.
- Randolph-Seng, Brandon, and Michael E. Nielsen. 2007. "Honesty: One Effect of Primed Religious Representations." *International Journal for the Psychology of Religion* 17: 303-315.
- Reicher, Stephen, and Mark Levine. 1994. "Deindividuation, Power Relations Between Groups and the Expression of Social Identity: The Effects of Visibility to the Out-Group." *British Journal of Social Psychology* 33(2): 145-163.
- Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. 2009. "Procrastination and Impatience." Northwestern University working paper.
- Rogers, William. 1993. "Regression Standard Errors in Clustered Samples." *Stata Technical Bulletin* 13: 19-23.
- Ruffle, Bradley J., and Richard Sosis. 2007. "Does it Pay to Pray? Costly Ritual and Cooperation." *The B.E. Journal of Economic Analysis & Policy* 7(1) (Contributions), Article 18.
- Sahm, Claudia R. 2007. "Stability of Risk Preference." Federal Reserve Board Finance and Economics Discussion Series 2007-66.
- Shariff, Azim F., and Ara Norenzayan. 2007. "God Is Watching You: Priming God Concepts Increases Prosocial Behavior in an Anonymous Economic Game." *Psychological Science* 18: 803-809.

- Shapiro, Jesse M. 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics* 89(2-3): 303-325.
- Stewart, Mark B. 1983. "On Least Squares Estimation when the Dependent Variable is Grouped." *Review of Economic Studies* 50(4): 737-753.
- Toburen, Tina, and Brian P. Meier. 2010. "Priming God-Related Concepts Increases Anxiety and Task Persistence." *Journal of Social and Clinical Psychology* 29(2), 127-143.
- Turner, John C. 1985. "Social Categorization and the Self-Concept: A Social Cognitive Theory of Group Behavior." In *Advances in Group Processes*, Vol. 2, ed. E. J. Lawler, 77-121. Greenwich, CT: JAI Press.
- Weber, Max, 1930. *The Protestant Ethic and the Spirit of Capitalism*. London: Allen and Unwin.
- Wheeler, S. Christian, and Richard E. Petty. 2001. "The Effects of Stereotype Activation on Behavior: A Review of Possible Mechanisms." *Psychological Bulletin* 127(6), 797-826.

Table 1. Public Goods Game Results

	Protestant	Catholic	Jewish	Agnostic/Atheist
Panel A. Amount contributed to public good				
<i>Religion salient</i>	0.15*	-0.18**	0.01	0.12
	(0.06)	(0.07)	(0.12)	(0.06)
Constant	0.53**	0.70**	0.56**	0.49**
	(0.05)	(0.05)	(0.08)	(0.04)
<i>N</i>	180	138	56	168
Panel B. Expectation of others' contribution to public good				
<i>Religion salient</i>	0.03	-0.12*	0.08	0.09
	(0.05)	(0.05)	(0.08)	(0.05)
Constant	0.59**	0.73**	0.61**	0.52**
	(0.03)	(0.04)	(0.05)	(0.03)
<i>N</i>	180	138	56	168
Panel C. Relationship between own contribution and expectation of others' contribution				
<i>Religion salient</i>	0.12**	-0.07	-0.07	0.04
	(0.04)	(0.05)	(0.09)	(0.05)
<i>E(Others' contribution)</i>	0.90**	0.94**	1.01**	0.95**
	(0.05)	(0.06)	(0.11)	(0.07)
Constant	-0.00	0.01	-0.06	-0.01
	(0.04)	(0.06)	(0.06)	(0.04)
<i>N</i>	180	138	56	168

Note: This table shows regression results where the dependent variable is the amount contributed to the public good (Panels A and C) or the expectation of others' average contribution to the public good (Panel B). *Religion salient* is a dummy for being in the religion-salient condition. *E(Others' contribution)* is the subject's expectation of other group members' average contribution. Huber-White standard errors are in parentheses below the point estimates.

* Significant at the 5% level. ** Significant at the 1% level.

Table 2. Risk Aversion Results

	Protestant	Catholic	Jewish	Agnostic/ Atheist
<i>Religion salient</i>	-0.03 (0.06)	-0.11 (0.07)	0.02 (0.06)	-0.12* (0.05)
<i>Larger Stakes</i>	0.27** (0.04)	0.31** (0.05)	0.26** (0.07)	0.26** (0.04)
Constant	0.23** (0.05)	0.21** (0.05)	0.09* (0.04)	0.16** (0.03)
$\hat{\sigma}$	0.39 (0.03)	0.36 (0.03)	0.26 (0.03)	0.32 (0.02)
<i>N</i>	242	154	80	196

Note: This table presents interval regressions where the latent dependent variable is the minimum risk premium required for a subject to accept a gamble. We pool each subject's two risk choices together. *Religion salient* is a dummy for being in the religion-salient condition. *Larger stakes* is a dummy for if the sure payout in the risky choice was \$100. The estimated conditional standard deviation of the latent dependent variable is denoted by $\hat{\sigma}$. Huber-White standard errors, clustered by subject, are reported in parentheses below the point estimates. The final row shows the number of reservation risk premium intervals in the regressions. * Significant at the 5 percent level. ** Significant at the 1 percent level.

Table 3. Discount Rate Results

	Protestant	Catholic	Jewish	Agnostic/Atheist
<i>Religion salient</i>	0.21 (0.36)	0.43 (0.50)	0.33 (0.57)	0.08 (0.51)
<i>1 week vs. 2 weeks</i>	0.01 (0.11)	-0.17 (0.17)	-0.37 (0.27)	-0.30 (0.17)
Constant	-2.96** (0.26)	-3.53** (0.36)	-3.62** (0.44)	-4.40** (0.42)
$\hat{\sigma}$	1.99 (0.16)	2.18 (0.21)	1.85 (0.20)	2.36 (0.18)
<i>N</i>	242	154	80	196

Note: This table presents interval regressions where the latent dependent variable is the log continuously compounded interest rate required to defer payment receipt. We pool each subject's two intertemporal choices together. *Religion salient* is a dummy for being in the religion-salient condition. *1 week vs. 2 weeks* is a dummy for the intertemporal choice being between payments deferred for one week versus two weeks. The estimated conditional standard deviation of the latent dependent variable is denoted by $\hat{\sigma}$. Huber-White standard errors, clustered by subject, are reported in parentheses below the point estimates. The final row of each panel reports the number of discount rate intervals in the regressions. * Significant at the 5 percent level. ** Significant at the 1 percent level.

Table 4. Dictator Game Results

	Protestant	Catholic	Jewish	Agnostic/Atheist
<i>Religion salient</i>	-0.03 (0.04)	-0.04 (0.04)	0.01 (0.07)	-0.05 (0.03)
Constant	0.22** (0.03)	0.17** (0.04)	0.14** (0.04)	0.17** (0.02)
<i>N</i>	186	139	62	176

Note: This table shows regression results where the dependent variable is the amount given away. *Religion salient* is a dummy for being in the religion-salient condition. Huber-White standard errors are in parentheses below the point estimates. ** Significant at the 1 percent level.

Table 5. Number of Anagrams Attempted

	Protestant	Catholic	Jewish	Agnostic/Atheist
<i>Religion salient</i>	-1.48 (2.53)	1.76 (3.98)	4.92 (5.99)	2.12 (2.63)
Constant	29.80** (1.89)	31.09** (2.89)	33.65** (4.48)	32.02** (1.78)
<i>N</i>	78	60	33	93

Note: This table shows regression results where the dependent variable is the number of anagrams attempted in the anagram-solving task. *Religion salient* is a dummy for being in the religion-salient condition. Huber-White standard errors are in parentheses below the point estimates. ** Significant at the 1% level.

Table 6. Bilateral Gift-Exchange Game Results

	Protestant	Catholic	Jewish	Agnostic/Atheist
Panel A. Slope from regression of work cost on wages between \$0 and \$3.50				
<i>Religion salient</i>	0.01 (0.02)	0.01 (0.02)	0.06* (0.03)	-0.02 (0.02)
Constant	0.08** (0.01)	0.08** (0.01)	0.05** (0.02)	0.09** (0.01)
<i>N</i>	104	91	40	113
Panel B. Wage offered as manager				
<i>Religion salient</i>	0.02 (0.19)	0.01 (0.19)	0.42 (0.29)	-0.28 (0.18)
Constant	1.41** (0.13)	1.28** (0.14)	0.98** (0.21)	1.57** (0.12)
<i>N</i>	104	91	40	113

Note: This table shows regression results where the dependent variable is the subject-specific slope coefficient from a regression of work cost expended as an Employee on managerial wage offered (Panel A) or the wage offered as a Manager (Panel B). *Religion salient* is a dummy for being in the religion-salient condition. Huber-White standard errors are in parentheses below the point estimates. ** Significant at the 1% level.

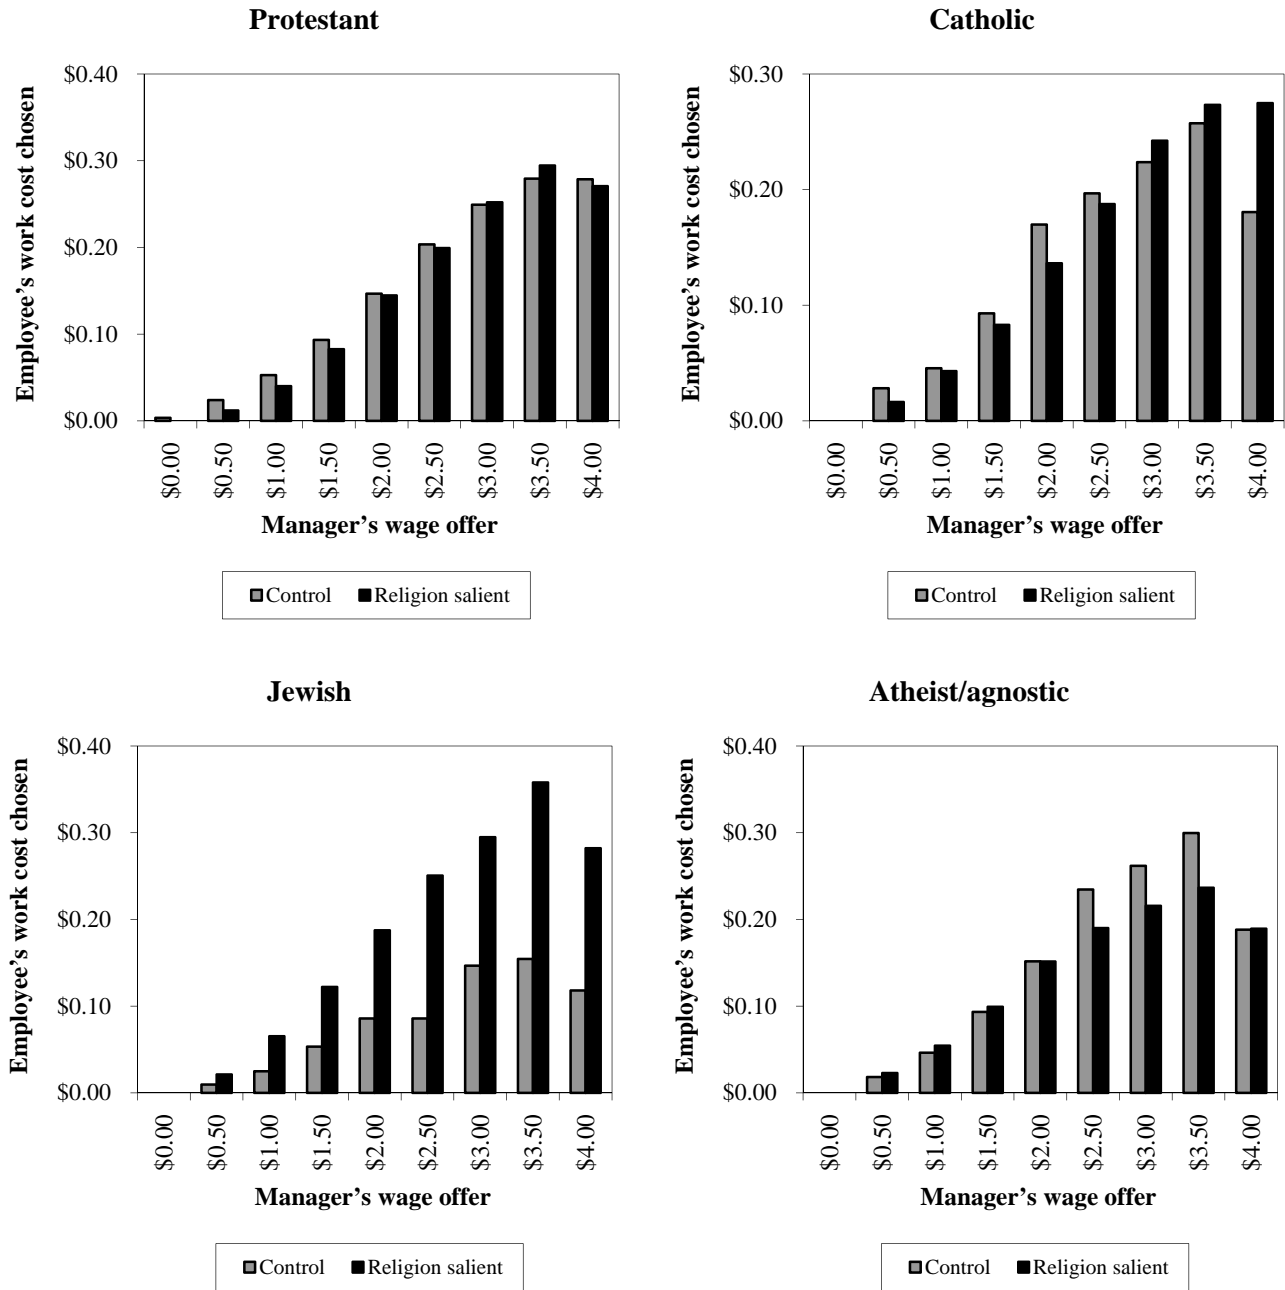


Figure 1. Average Employee work costs chosen in response to managerial wage offers in gift exchange game.